Silicon Graphics, Inc.

SGI's First Approach to Multi-paradigm Computing

Bill Mannel
Director of Marketing
Server & Platform Group

sgi®

SGI Scalable ccNUMA Architecture

Basic Node Structure and Interconnect

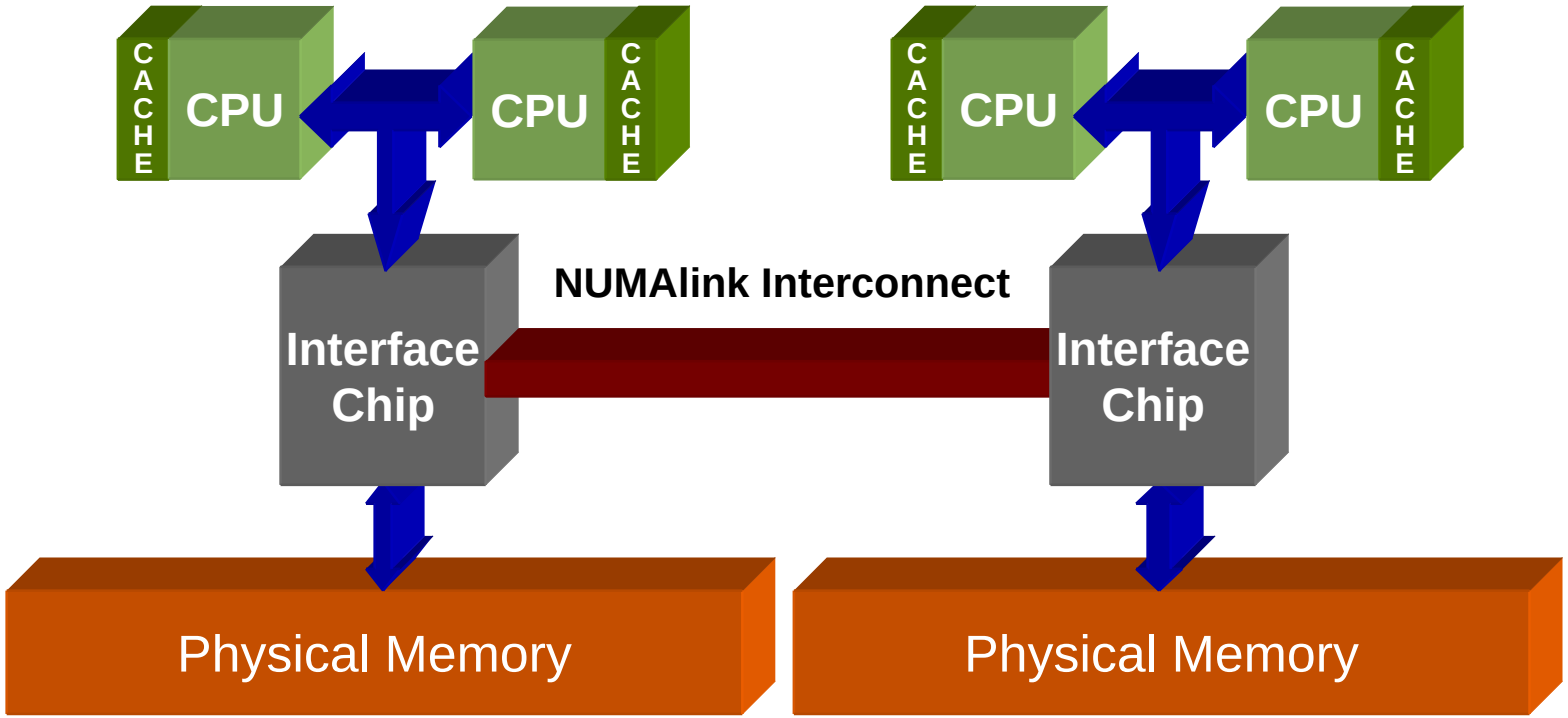

SGI Scalable ccNUMA Architecture

Basic Node Structure and Interconnect

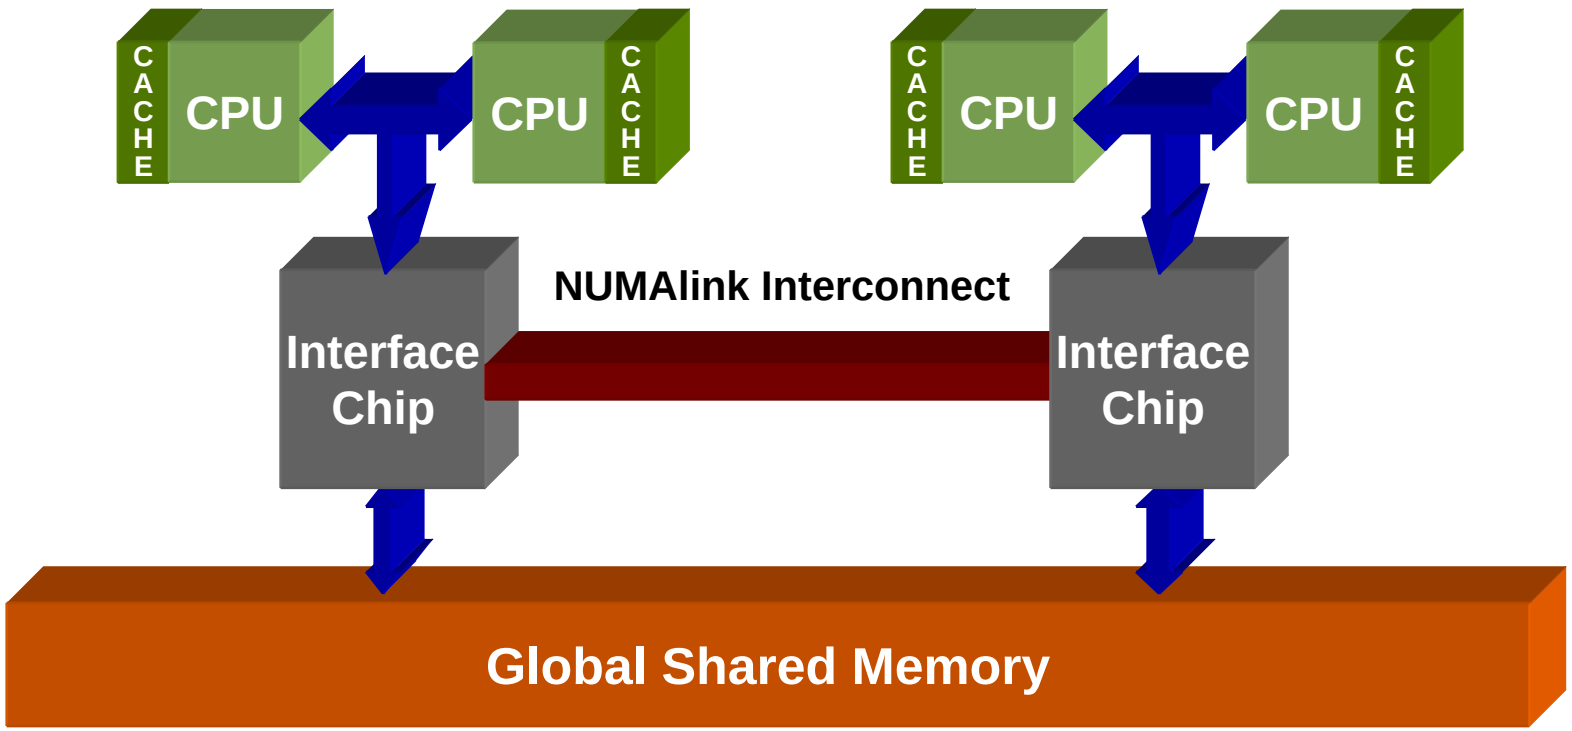

SGI Scalable ccNUMA Architecture

Scaling to Large Node Counts

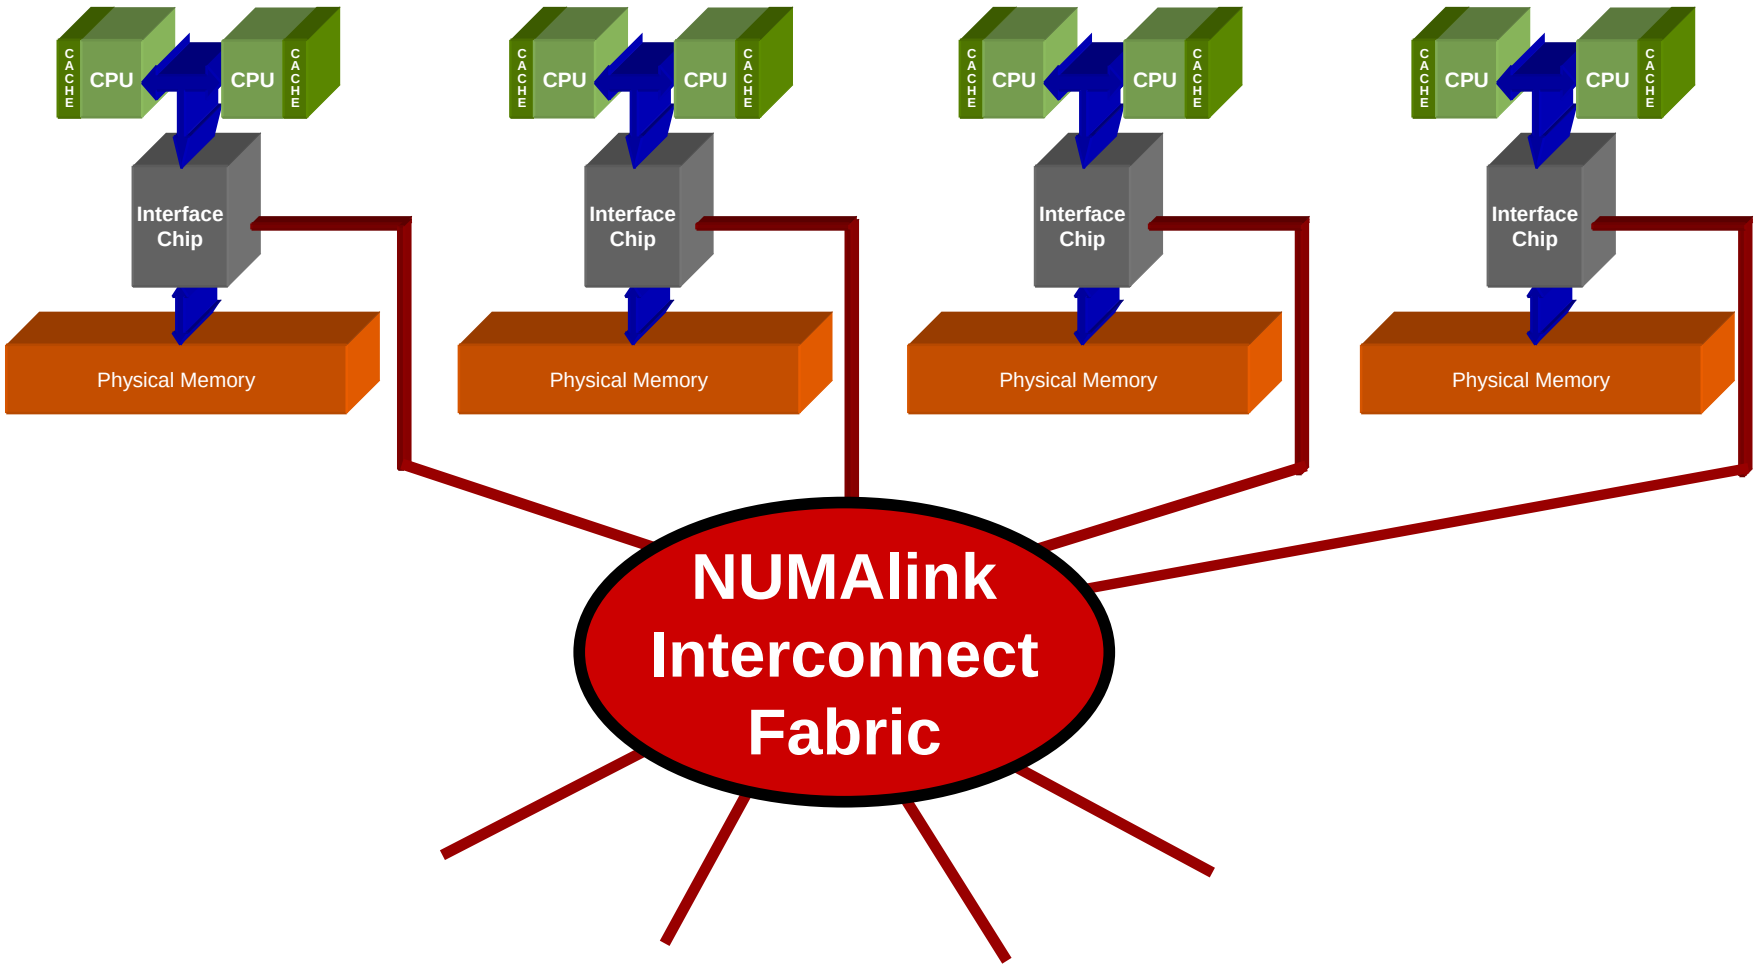

SGI Scalable ccNUMA Architecture

Multi-Paradigm Computing Architecture

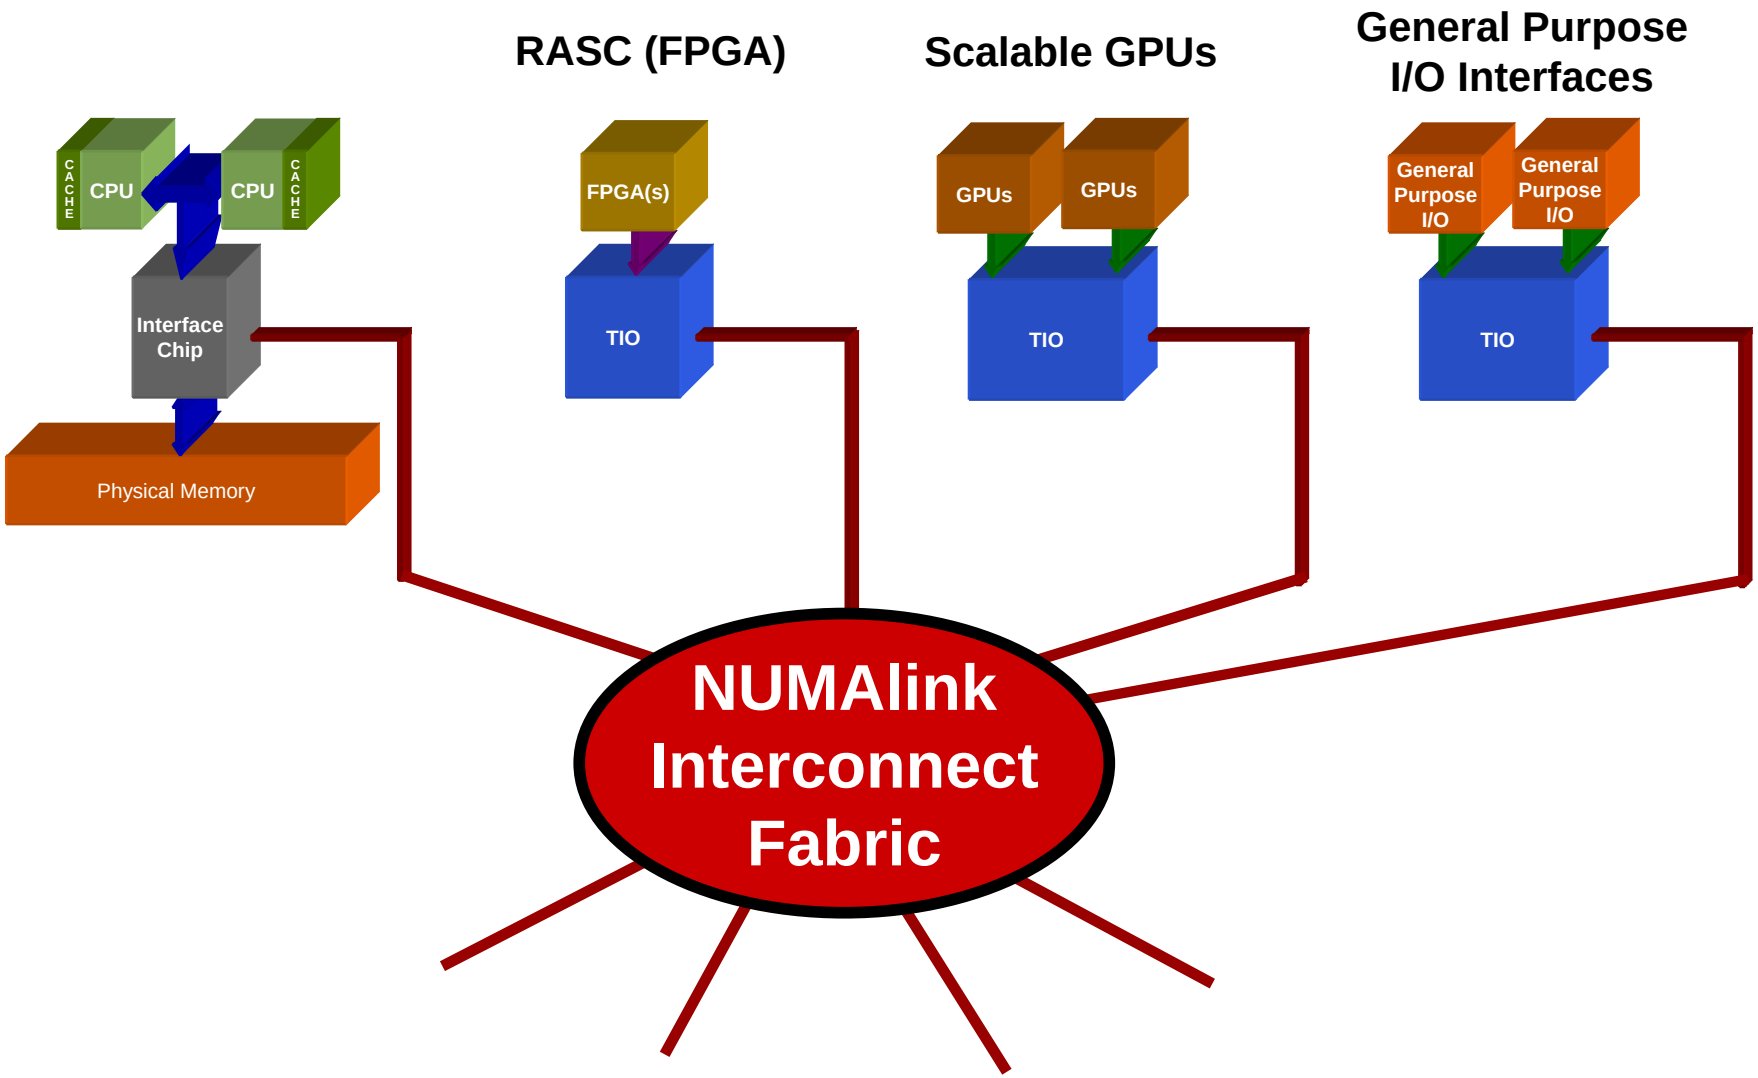

RASC

- RASC is *Reconfigurable* Application Specific Computing
- First instantiation of RASC is FPGA's, *Field Programmable Gate Arrays*

Why FPGAs…?

In a word…performance

- Some apps…orders of magnitude perf. increases
- Low power
- Broader market use ensures that technology will progress from 90nm to 65nm and beyond…

HPC Use of FPGAs…Challenges

- **Performance**
  - Bandwidth to/from system
  - Scalability
- **Ease of Use**
  - Languages
  - Compilers
  - Debuggers
  - APIs

Addressing the Performance Challenge
Integration into NUMAlink

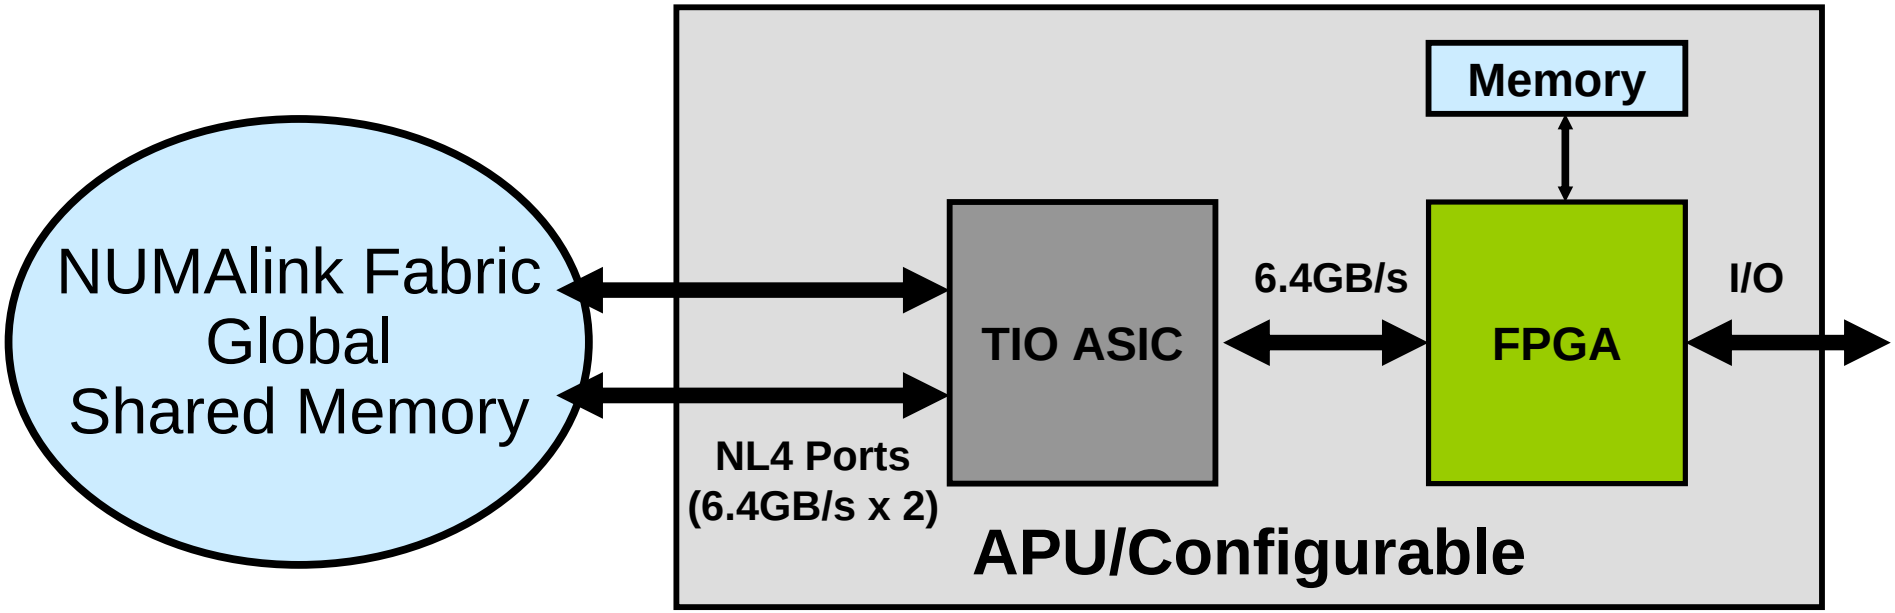

- Direct connect to Numalink (6.4 Gb/sec)
- Fast System Level Reprogramming of FPGA
- Atomic Memory Operations
- Hardware Barriers—load balancing
- Configurations to 128 NUMA/FPGA Nodes

MOATB--Early Access System Configuration

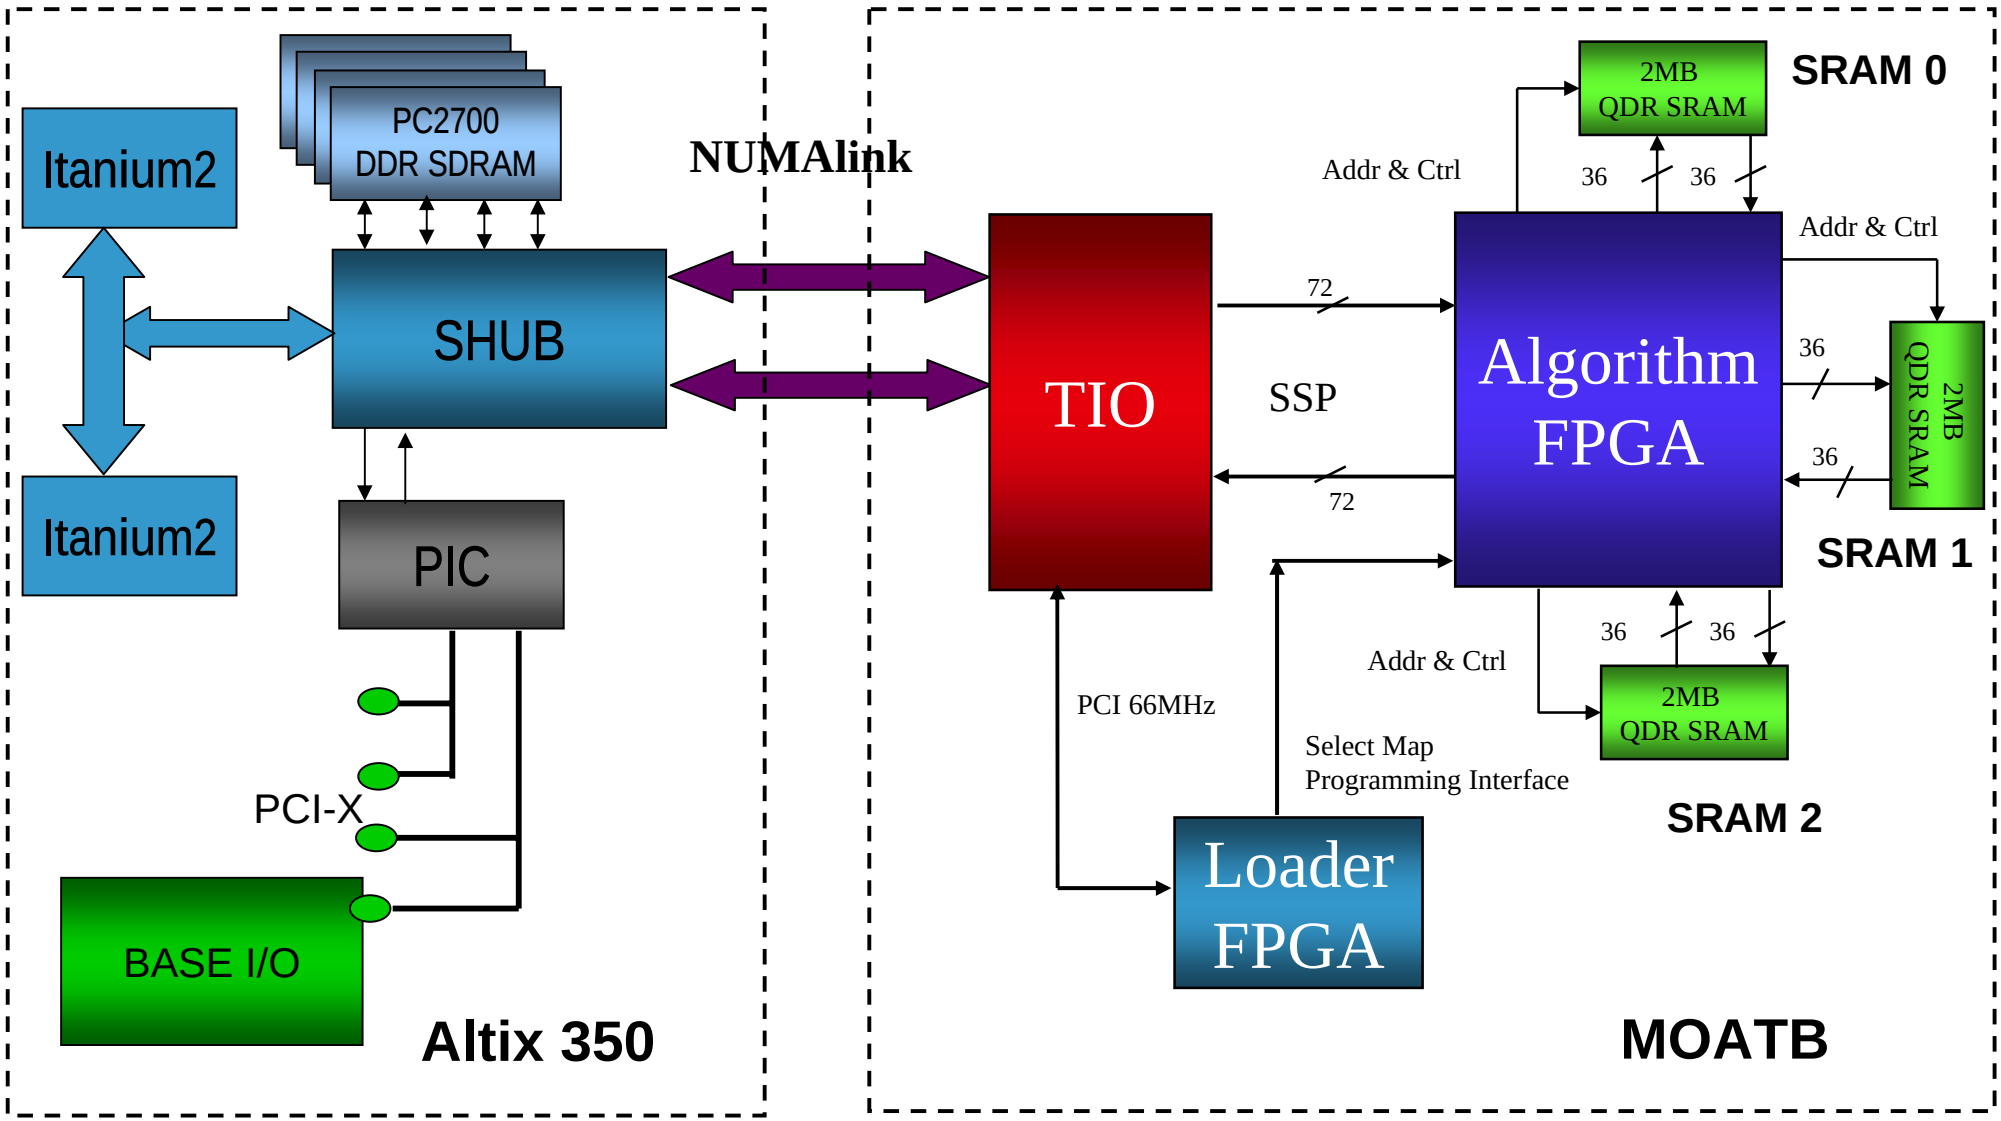

MOATB Early Access System

Current Customers

- NSA (2 units)
- U of FL
- NCSA (2 units)
- **NRL/US Navy**
- GWU
- PNNL (2 units)
- DSSI
- Starbridge Systems

Add'l Interest

- NexQL—Athena
- LANL
- Northrop (Azusa)
- Queens Univ. Belfast—2 x Athena
- GMU--Athena

sgi

MOATB Sample Application Performance

Bit Manipulation (Crypto)

79x 1.5GHz Itanium-2 (single MOATB)

119x 1.5GHz Itanium-2 (dual MOATB)

Graphics Edge Detection

42x 1.5GHz Itanium-2 (single MOATB)…demo'd at NAB

Customer Application

105 speedup on scalar microprocessor**

Next Product - Athena Computation Brick, summer 2005, Cooperation with Nallatech & SGI

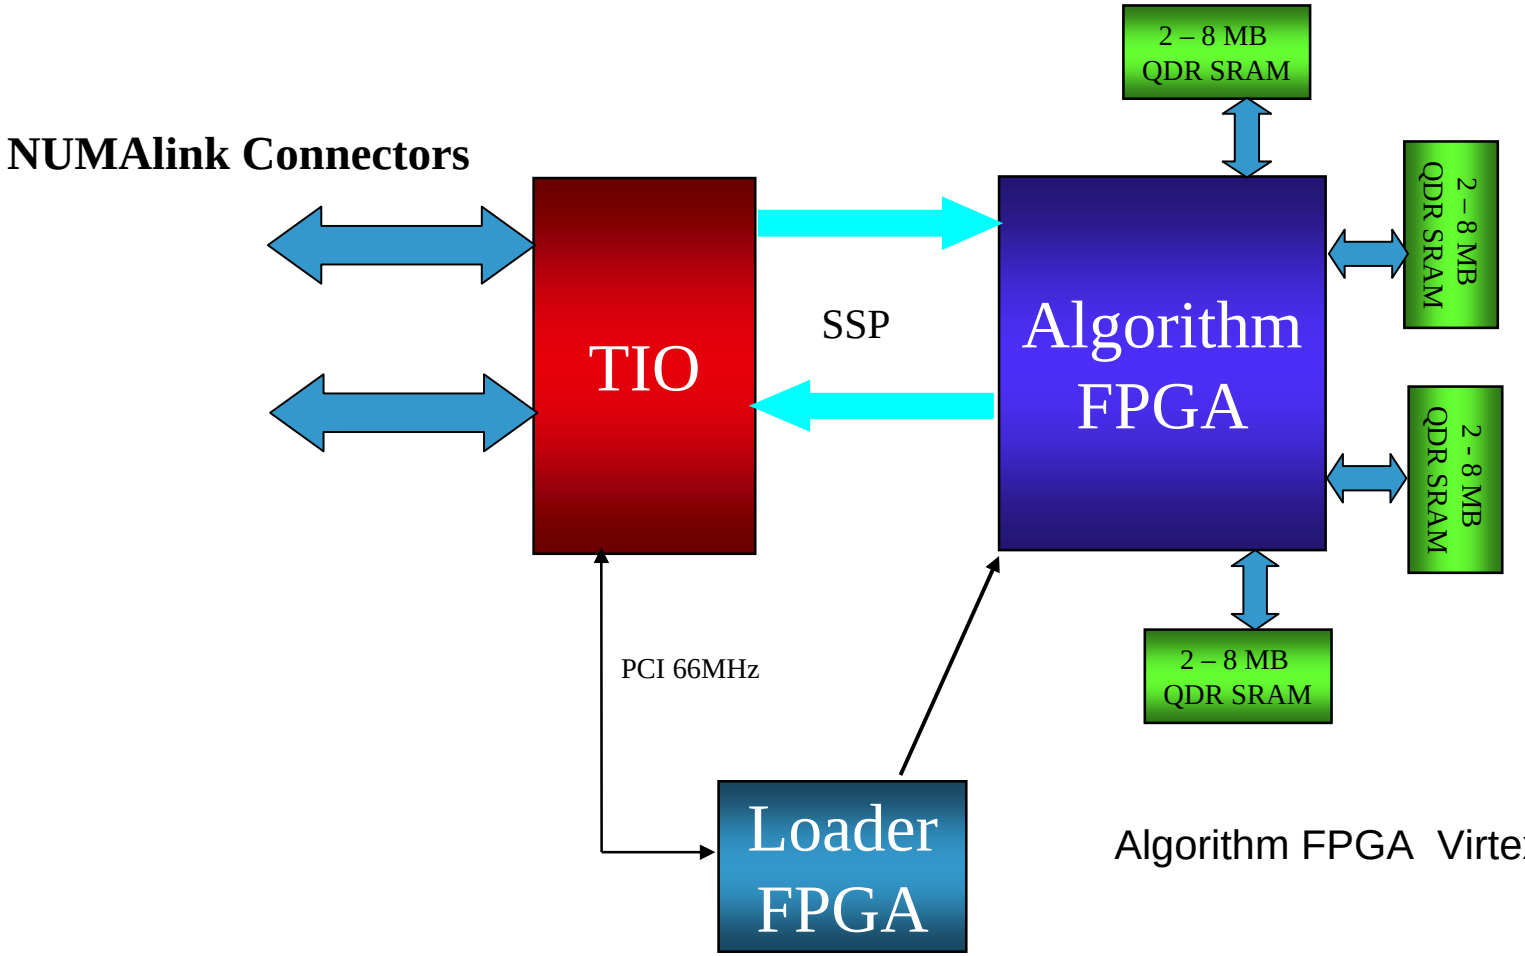

Algorithm FPGA Virtex2 6000 -6

Abacus Computation *Blade*

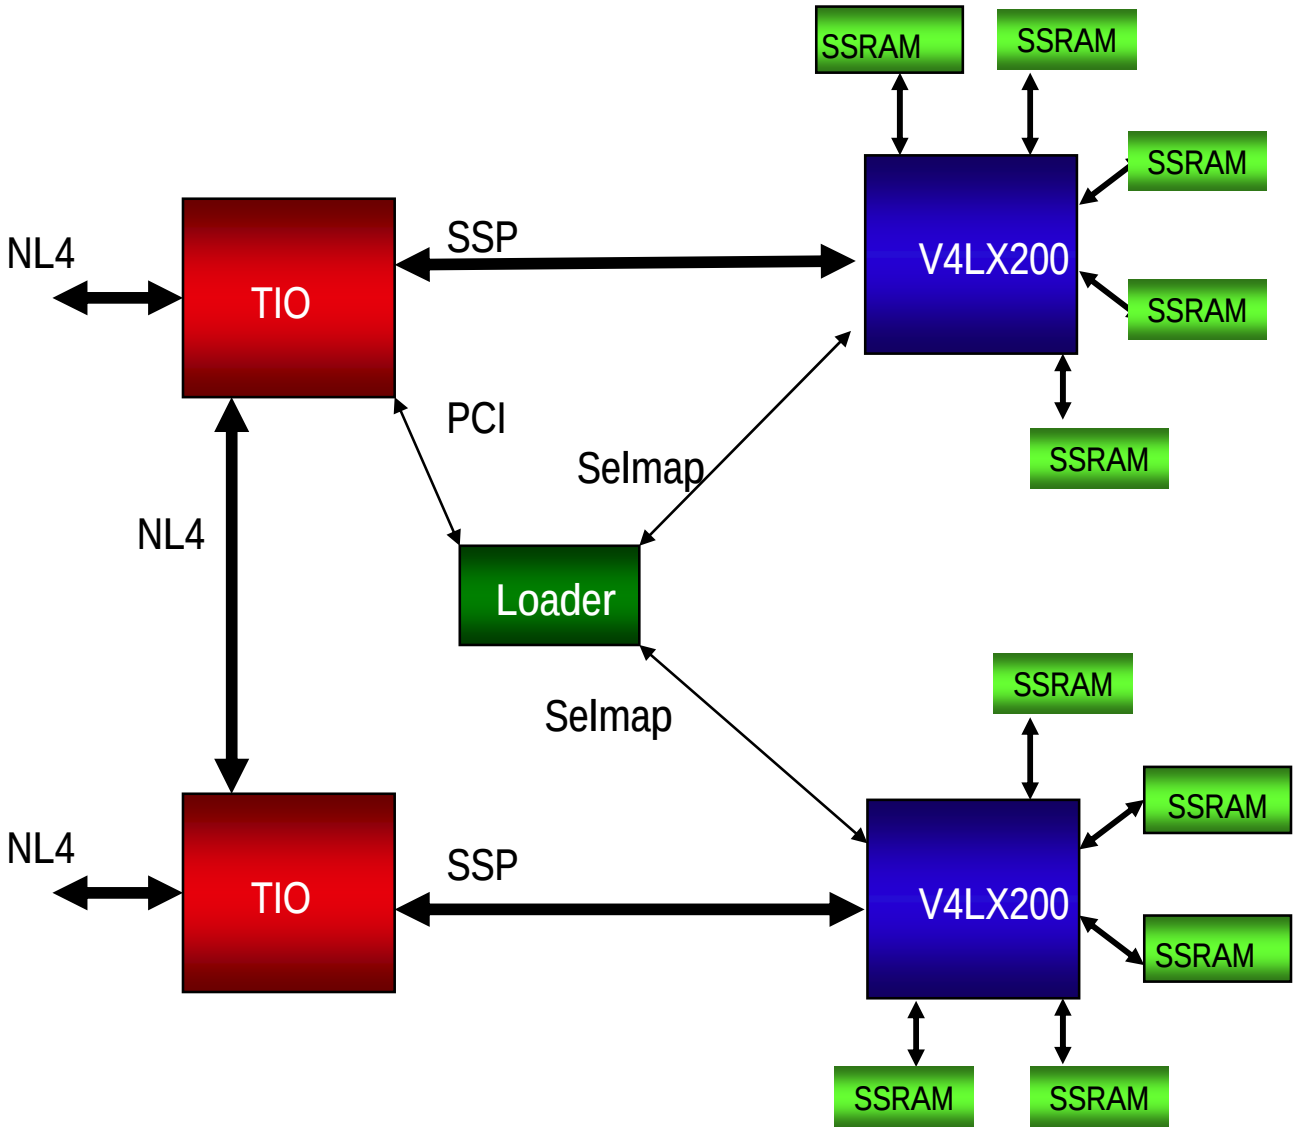

Abacus Layout

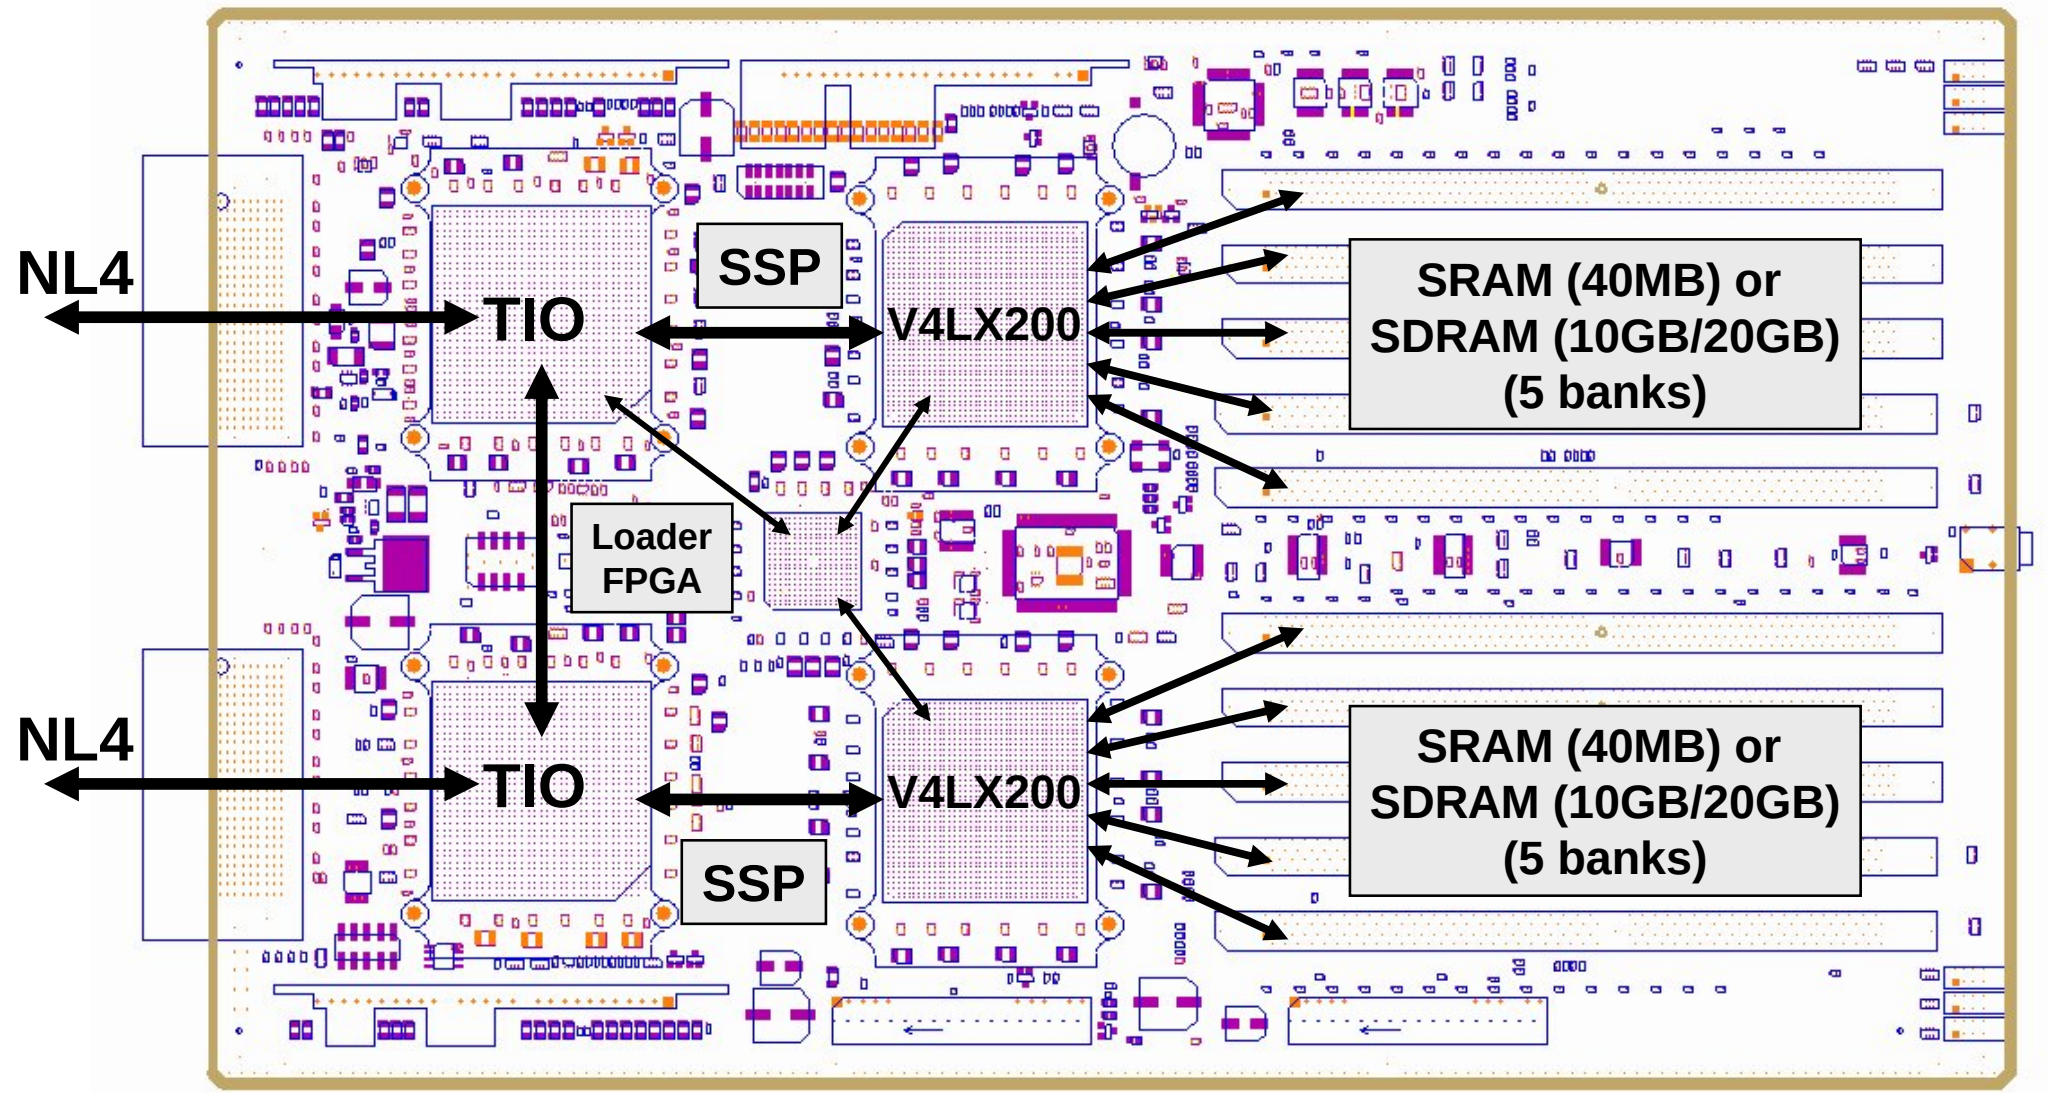

Proposed Multi-FPGA *Blade*, early CY06

TIO/SSP + FPGA modules

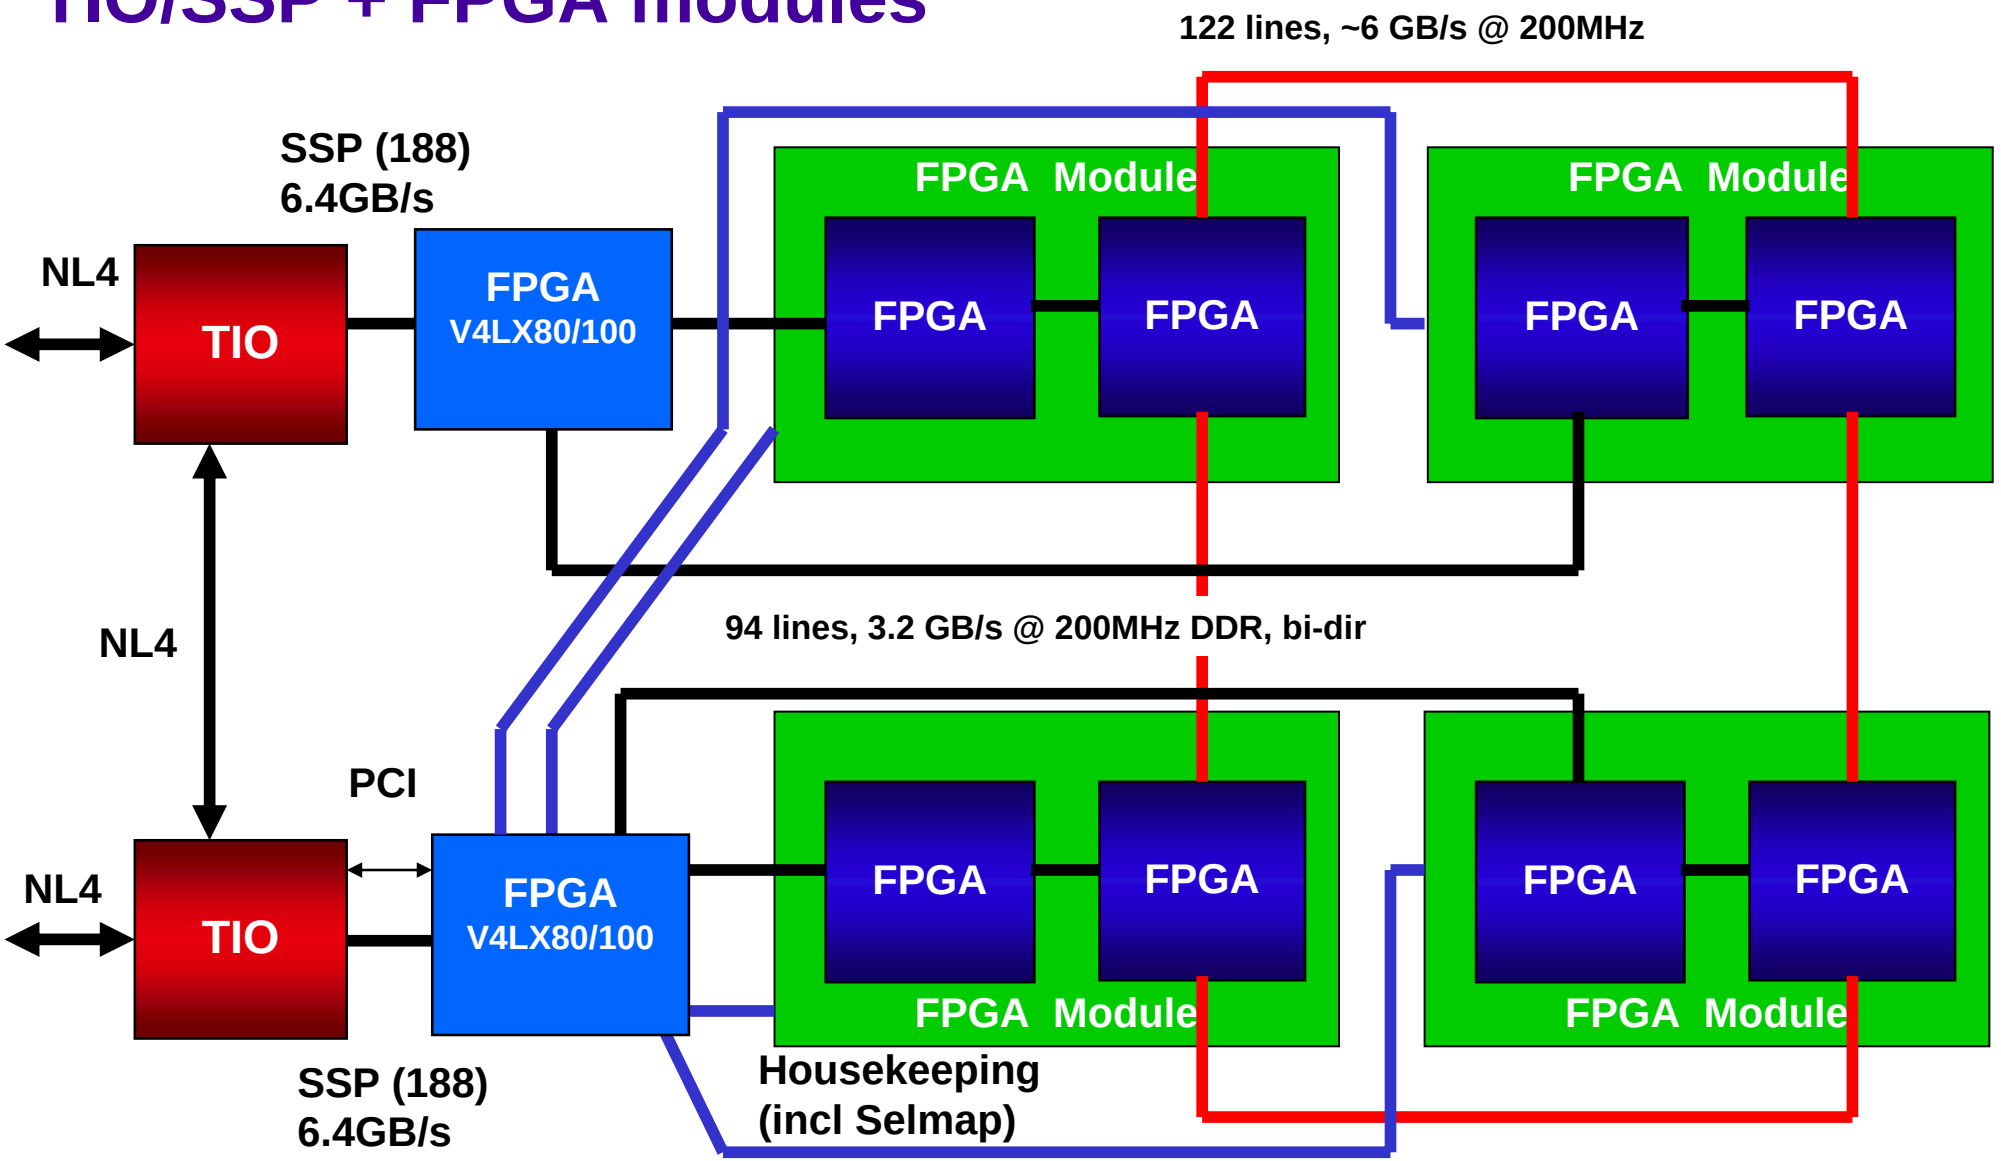

Development Environment
RASC Software Stack

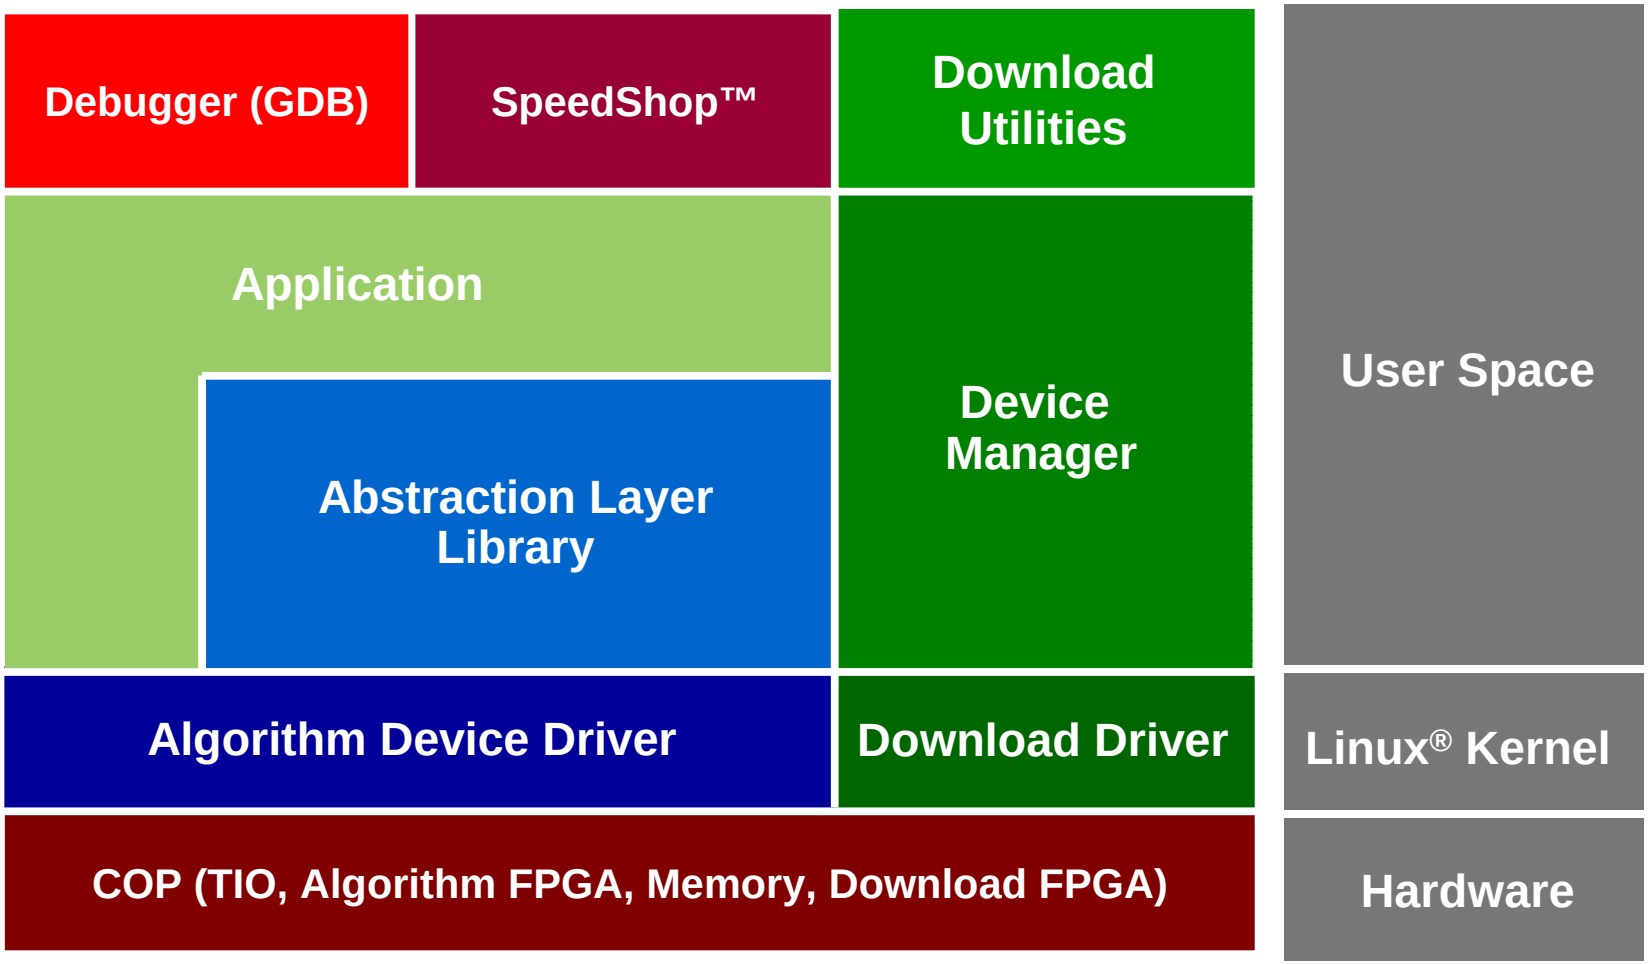

High Level Languages (HLL)

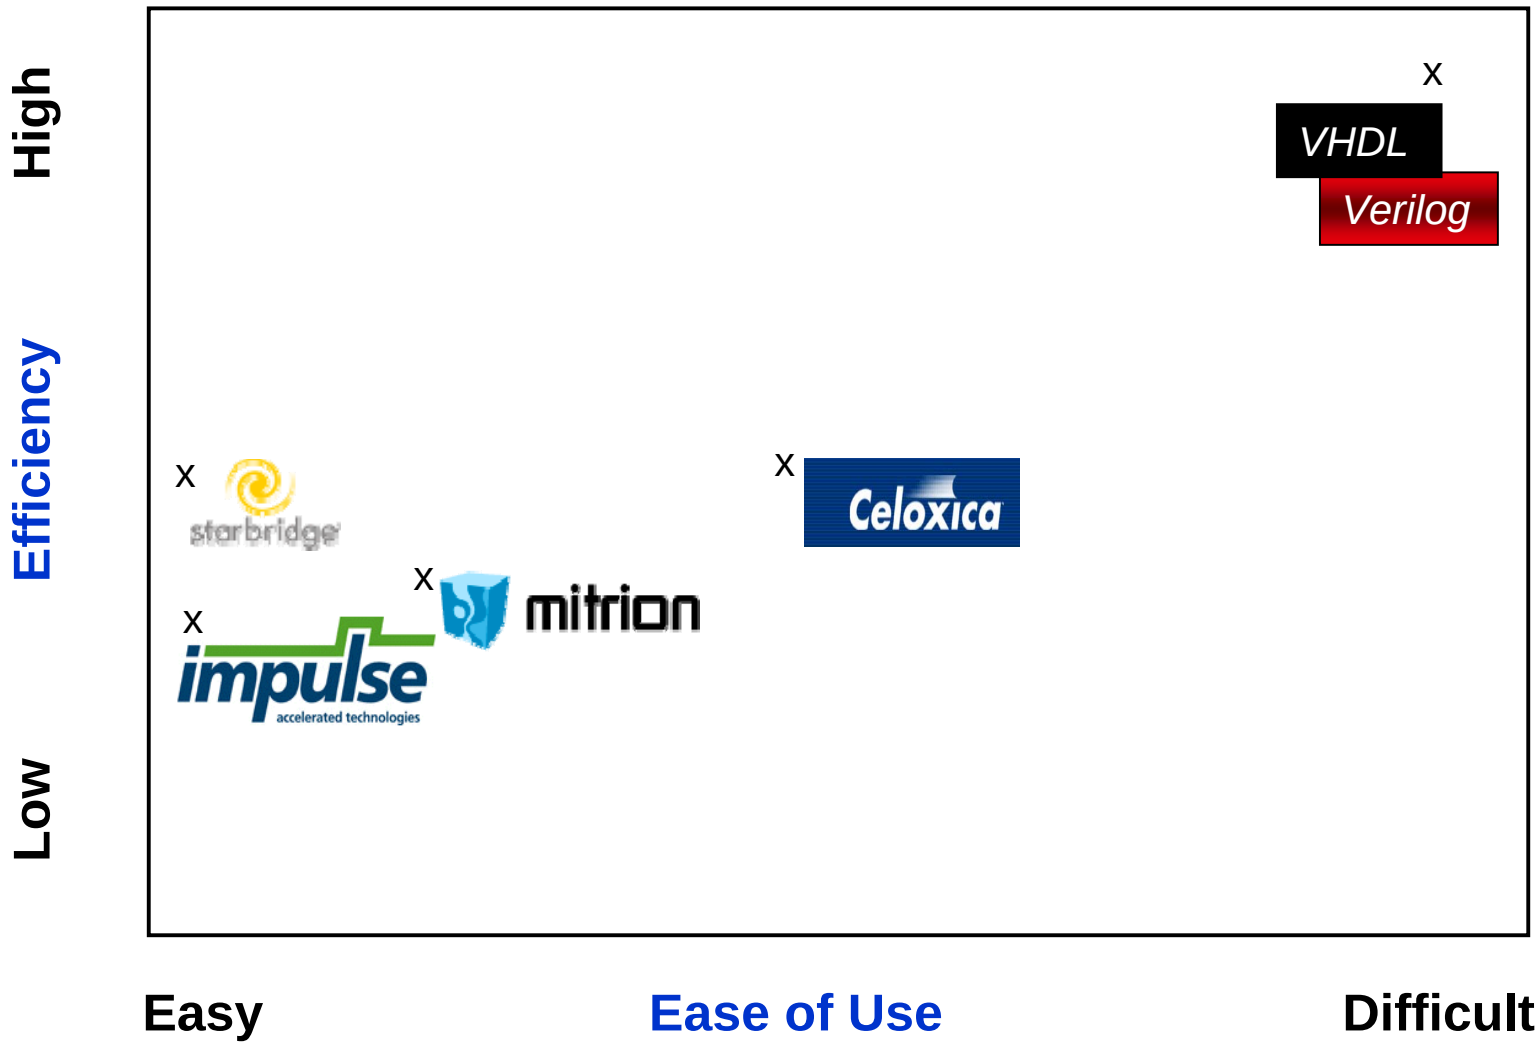

FPGA Aware Debugger

- Based on Open Source Gnu Debugger (GDB)
- Uses extensions to current command set
- Can debug host application and FPGA
- Provides notification when FPGA starts or stops
- Supplies information on FPGA characteristics
- Can “single-step” or “run N steps” of the algorithm
- Can HLL line step per C-line source
- Dumps data regarding the set of “registers” that are visible when the FPGA is active

Abstraction Layer: Algorithm API

The Abstraction Layer's algorithm API mirrors the COP API with a few additions that enable:

Wide Scaling

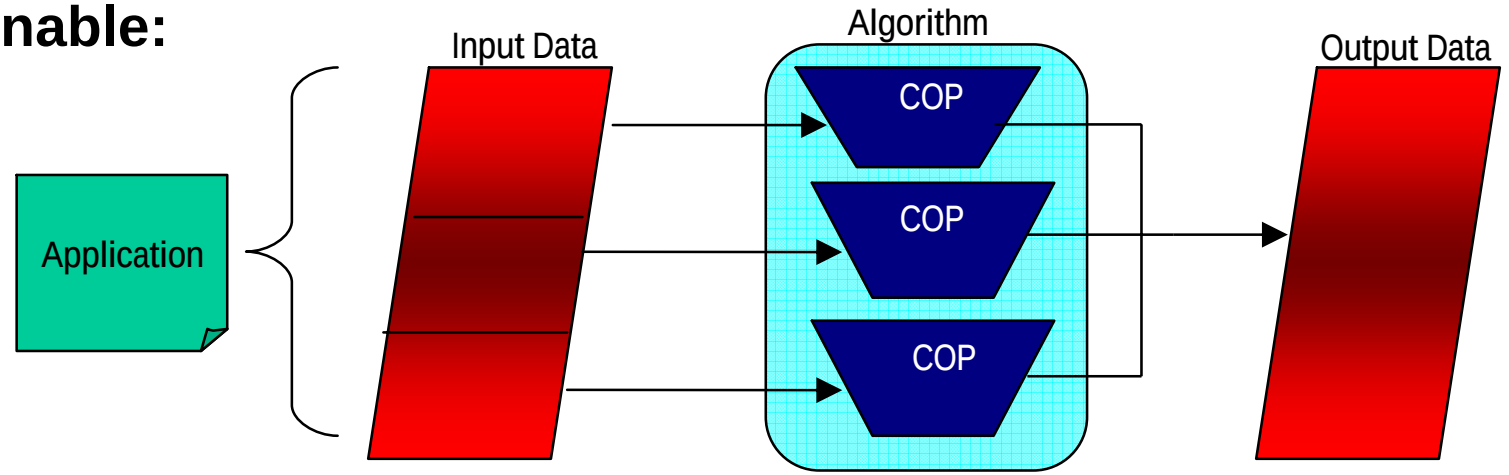

- and -

Deep Scaling

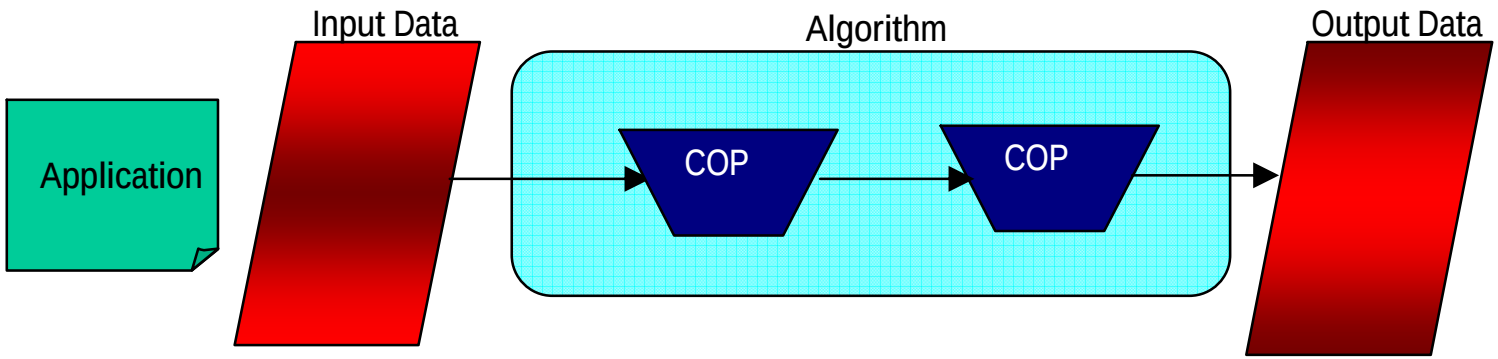

Working with industry/customers (openfpga.org) on API stds…

sgi

Verilog / VHDL Module Support

- Templates for Verilog and VHDL
  - Fast start to algorithm coding
- Provide a system simulation stub
  - Allows both simulation debug or system debug
- Provide source code for core service
  - Allows user to modify to meet special needs
- Extractor tools supports GDB meta-data
  - Application and FPGA debugging

Multi-Paradigm Computing
Other Non-traditional Processing Initiatives

- GPU-based processing
  - High potential performance (200-300GF peak today) and performance/price on single precision floating point applications…clear roadmap to future semiconductor process technologies
  - SGI working with SI on scaling to multiple GPUs and on development environment/programming paradigms…initial focus on signal processing apps
- Specialized processors…Clearspeed processors, custom processors (MD-GRAPE, classified chip)
  - High potential performance/watt on certain apps
  - Not clear that market exists which would fund semiconductor process migration roadmap…?

Multi-Paradigm Computing

Terascale to Petascale Data Set...Bring Function to Data

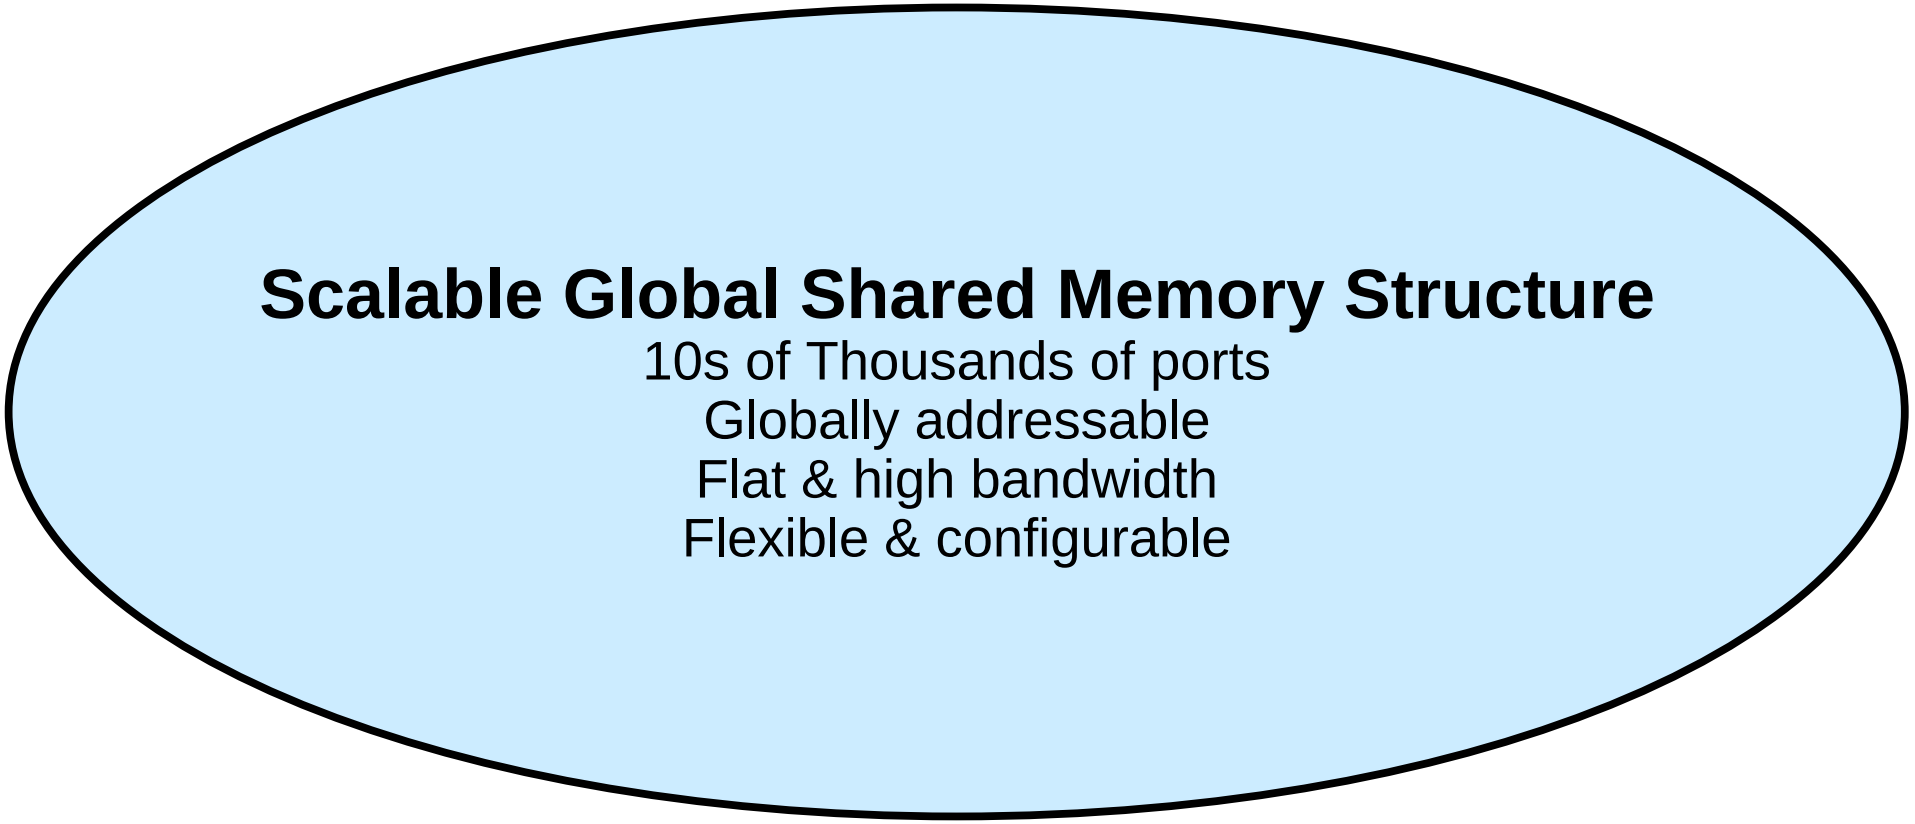

- **Globally addressable, universally accessible, high-bandwidth flat address space**
- **Addressing the transition from Terascale to Petascale computing**
- **Tightly coupling computation, I/O, and visualization functions to memory**

sgi

Multi-Paradigm Computing

Terascale to Petascale Data Set...Bring Function to Data

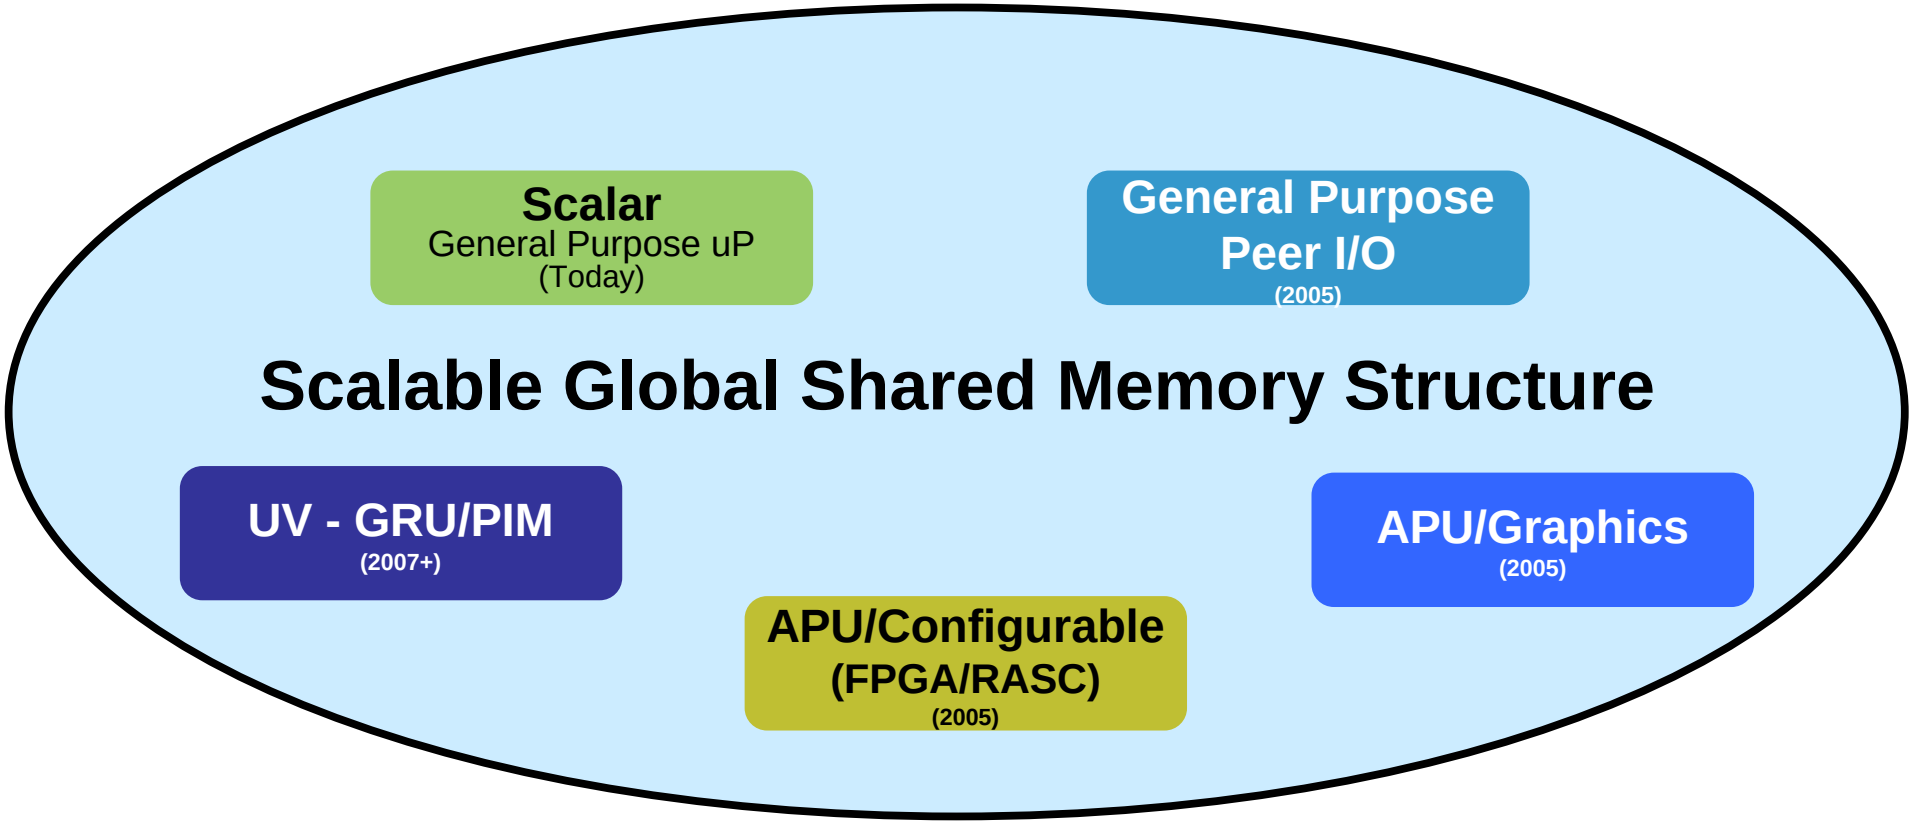

- Globally addressable, universally accessible, high-bandwidth flat address space
- Addressing the transition from Terascale to Petascale computing
- Tightly coupling computation, I/O, and visualization functions to memory

sgi